REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	30 May 2003	FINAL 23 March 2001 - 30 May 2003

4. TITLE AND SUBTITLE
"COMPUTER ASSISTED REASONING FOR ROBUST STRATEGIES: SCIENTIFIC AND TECHNICAL FINAL REPORT"

5. FUNDING NUMBERS
C DAAH01-01-C-R161

ARPA Order K724/26

6. AUTHOR(S)
Steven C. Bankes, Robert J. Lempert, Steven W. Popper

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Evolving Logic
21280 Coline Dr.
Topanga, CA 90290

8. PERFORMING ORGANIZATION REPORT NUMBER
A004

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency (DOD)
3701 North Fairfax Dr.
Arlington, VA 22203-1714

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the authors and should not be construed as an official DARPA or US government position, policy or decision, unless so designated by other documentation.

12 a. DISTRIBUTION / AVAILABILITY STATEMENT	12 b. DISTRIBUTION CODE
Approved for public release; distribution unlimited.	

13. ABSTRACT (Maximum 200 words)

1

Evolving Logic successfully completed a project to build upon its existing, patented Computer Assisted Reasoning® system (CARs™) technology. CARs 2.0 allows users to generate compound computational experiments as well as to analyze their outputs. This is key to creating a capability that for the first time allows users to reason under conditions of complexity and deep uncertainty in order to develop and test robust adaptive strategies.

This capability is essential to gaining operational facility in asymmetric warfighting environments, for assessing threats, and developing postures that are robust to surprise. The document details technical capabilities within CARs™ and reports extensively on feedback from various user communities.

14. SUBJECT TERMS
Strategic planning; Robust Decisions; Adversarial reasoning; Risk Assessment; Surprise management; Model-based planning; Decision-support software; Scenario planning;

15. NUMBER OF PAGES
37

16. PRICE CODE

17. SECURITY CLASSIFICATION OR REPORT	18. SECURITY CLASSIFICATION ON THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev.2-89)
Prescribed by ANSI Std. 239-18
298-102

Enclosure 1

20030801 132

SCIENTIFIC AND TECHNICAL FINAL REPORT:

"COMPUTER ASSISTED REASONING FOR ROBUST STRATEGIES"

30 May 2003

Sponsored by

Defense Advanced Research Projects Agency (DOD)

Information Awareness Office

ARPA Order K724/26

Issued by U.S. Army Aviation and Missile Command Under

Contract No. DAAH01-01-C-R161

Contractor: Evolving Logic
21280 Colina Dr.
Topanga, CA 90290

Principal Investigator: Dr. Steven C. Bankes
(310) 836-0958

Effective Date of Contract: 22 March 2001
Contract End Date: 22 June 2003

Short Title:
"Computer Assisted Robust Strategies"

Report Period: 23 March 2001 to 30 May 2003

ABSTRACT: Evolving Logic successfully completed a project to build upon its existing, patented Computer Assisted Reasoning® system (CARs™) technology. CARs 2.0 allows users to generate compound computational experiments as well as to analyze their outputs. This is key to creating a capability that for the first time allows users to reason under conditions of complexity and deep uncertainty in order to develop and test robust adaptive strategies. This capability is essential to gaining operational facility in asymmetric warfighting environments, for assessing threats, and developing postures that are robust to surprise. The document details technical capabilities within CARs™ and reports extensively on feedback from various user communities.

I. INTRODUCTION

"Uncertainty -- in the economy, society, politics -- has become so great as to render futile, if not counterproductive, the kind of planning most companies still practice: forecasting based on probabilities"
Peter Drucker
Wall St. Journal, 22 July 1992

"The only surprise is that we are surprised when we are surprised."
Attributed to Donald Rumsfeld, Secretary of Defense

Decision-makers -- from military planners to managers of corporate supply chains -- often confront conditions of *deep uncertainty* where they do not know, or cannot agree upon, the probabilities to assign to alternative scenarios or even the best models to describe the future. Deep uncertainty is ubiquitous. It can arise in times of fast technological and societal change or when adversaries are actively trying to counter your actions. In such situations, standard decision support tools are not straightforward to apply and often will not accurately support the aims of decision-makers. Traditional decision support tools have too frequently failed to deliver on their promise because they do not address the deep uncertainty that decision-makers know they face.

Evolving Logic entered into this project possessing patent pending prototype technology designed to generate and display results from compound computational experiments. This technology had already been applied in the form of a software package known as CARs™ (Computer Assisted Reasoning® system) to the support of planning and decision making under conditions of deep uncertainty. To do so, CARs™ supports Evolving Logic's Robust Adaptive Planning™ robust decision method. This was before

September 11, 2001. Since that date it has become clear that in considering Concepts of Operations and Course of Action planning within the asymmetric warfare environment characterized by the post-September 11 world, precisely the types of capabilities provided by CARs are required to meet the challenges we now face.

The purpose of this Phase II project was to build a prototype decision support system for robust decision making under conditions of deep uncertainty. The project was designed to improve upon our existing commercial product to help individuals and groups use the full range of available information, from quantitative data to intuition, to identify robust strategies. The basic idea behind CARs™ is to use computer models to generate a large ensemble of plausible future scenarios, where each scenario represents one estimate of the way the world works (including potential models of adversarial reasoning) and one choice among alternative plans for acting on the world. We then use computer search and visualization methods to help decision-makers extract information from this ensemble of scenarios to evaluate alternative decisions. This present project thus had both technical goals as well as a need for considering issues of user interface and end-use.

The technical goals and milestones have been the subject of monthly reports since project inception. Therefore, they will be discussed below only in overview in the section following this one. Evolving Logic succeeded in building upon its existing CARs™ platfrom to produce a fully capable, general purpose system that can operate on a variety platforms (e.g. Windows, Linux); is configured for various modes of use (e.g. stand alone, multi-user, web deliverable); links easily to existing software packages for statistical analysis, database storage, simulation modeling, and graphical visualization; and can exploit peer-to-peer computational resource sharing. This involved a partial reimplemention of the existing CARs™ system, to expand its utility and improve on its design where experience has revealed improved design options. In particular, much C++ code was rewritten in Java in order to provide maximum platform independence while maintaining operating speed. The control flow and thread management architecture of the system was enhanced to enable greater flexibility in the ways users can interact with the system, including the deployment of the system in multi-user, cooperative planning modes.

The second main focus was to establish as a key to commercialization and broad military application, systematic procedures and the necessary software components for

rapid customization of CARs™ to particular applications. This involved developing procedures to elicit from user groups the best visualizations and sequences in which to present these visualizations in order to convey the information most crucial to informed, successful decision-making. It also meant developing a suite of what we call derivation methods for CARs™ to make it easy for users to construct the sometimes complex series of queries to the underlying models and data necessary to answer the particular questions of interest for their robust decision making application.

This second effort could only be pursued once the first was well in hand. In many ways, this is the crux of the matter – demonstrating that substantive user communities can interact with and utilize the output from CARs™ to achieve better means for making decisions under complexity and deep uncertainty. The third section of this report discusses in detail the user benchmarking and rapid prototyping efforts that were part of this project.

Finally, the last section lays out plans and prospects for Phase III commercialization of the results from the Phase II project.

II. OVERVIEW OF TECHNICAL RESULTS

All features of the previous system (CARs 1.0) have been re-implemented in a new version (called at differing points in time CARMEN or CARs 2.0) and a number of new features have been innovated. Our approach has been to utilize Evolving Logic's existing software base both as a starting point for creating a new system and as scaffolding for rapid prototyping new features. Design insights we have gleaned from using our previous technology (CARs™) have been incorporated into the new system. And a variety of desired new features have been implemented as will be described below. The system has been used on several internal projects at Evolving Logic, and prototype versions have been given, on an alpha-test basis, to colleagues at other institutions. Development also was aided by offering customized application versions of the system for sale to commercial clients. This work culminated in an alpha release of the new Windows version of the system to a limited community of users, and we continue to develop new capabilities.

Evolving Logic's foundation technology is an object oriented framework known as SPICE (Syntactic Protocol for Instantiating Computational Experiments). SPICE provides a

general-purpose engine for reasoning with computational experiments by offering a convenient platform for linking the numerous functions necessary to support these experiments and user interactions with them. The SPICE mechanism can encapsulate a wide variety of sources for computational experiments such as simulation models, statistical analyses, neural net training processes, and use large numbers of such experiments to support reasoning and problem solving. The technology allows us to capture both knowledge about and deep uncertainty regarding complex systems by representing infinite ensembles of alternative plausible models. SPICE allows reasoning about the properties of such ensembles, and hence the properties of ensembles of possible future events, through tailored constellations of computational experiments whose specification is guided by user goals and reasoning strategies.

Language. In our previous (CARs™ 1.0) technology, the SPICE mechanism was implemented in C++ (with a graphical user interface written in Java). We have re-implemented SPICE entirely in Java. In doing so we have made several improvements in the design of the SPICE base classes (affecting primarily the relation between Contexts, ResultsStores, and Derivations) over that in CAR™s 1.0. In addition to these improvements, the re-implementation greatly eases porting the system to alternative computers and operating systems. In particular, CARs™ has now been ported into the Windows operating system (specifically Windows 2000). Initial porting experiments to the Macintosh OS have also been undertaken.

Architecture. During this period we designed and implemented a multi-threaded process architecture, extending the existing SPICE base classes. This multi-threaded architecture will enhance CARs™ functionality significantly. Multi-threading allows us to experiment with regimes where user inputs, real time data, and emerging results from ongoing computational experiments can interact in diverse ways and will enable new options for user interaction with dynamic planning environments. It will also enable new approaches to experimental design where users interactively "fly" the search algorithms as they explore the scenario space. Such approaches may significantly improve users' ability to test hypotheses and intuitions against the information contained in high-dimensional models and large datasets. This enhancement will allow the system to operate in a multi-

user mode, allowing cooperative planning systems to be built, such as might be used in collaborative, web-based planning tools.

Interfaces. Our re-implementation of the SPICE base classes has been designed to facilitate easy interfaces with a wide variety of other software. These internal APIs will allow us to link with many types of computer simulation and data fitting models, generators such as search algorithms, existing software packages for database storage, statistical modeling, and graphical visualization.

The most obvious effect of these advances has been creating an explicit harness for running models written in Java. CARs 1.0 only supported the import of models through C++. Further, a mechanism has been designed for a harness that would use the Common Object Model (COM) interface under Windows. Now implemented, this allows us to drive models written in any modeling system that supports this interface, including Excel and Analytica.

Further, we have designed and implemented a new facility in our internal data flows, called "generalized bookmarks" which allow specific cases from one computational context to be used to create derived contexts, compound experiments, or specific visualizations in other contexts very flexibly and powerfully. This supports, for example, using searches to find extreme cases which automatically are used in lists of scenarios, allowing the user to graphically visualize outputs in one context that correspond to some interesting case in another context, and "drilling down" where by clinking on a point in a visualization, the user will be able to set off arbitrarily complex computations whose specifications involve the case corresponding to the point on the visualization selected.

Generators. The CARs™ approach to robust decision making requires a variety of generators to assist the user in choosing cases with which to populate the ensemble of scenarios. All generators that were supported in CARs™ 1.0 have been reimplemented in CARs™ 2.0 and we have identified the highest value algorithms from diverse fields such as operations research, machine learning, and statistical design of experiments, and have begun implementing them for use in our system.

Searches. In addition, we have developed means for implementing generators for uses that are relatively unprecedented in computational science. An important class of

generators for CARs™ are searches across the ensemble of scenarios. Searches are integral to implementing the RAP™ method for robust decision making for two reasons. First, our concept of robust strategies requires us to consider level sets of models, priors, and/or strategies that satisfy some conditions or constraints. Searches are integral to representing such sets on a computer. CARs™ now generates such sets by launching searches that generate a finite set of exemplar points or by finding surfaces in multi-dimensional space that separate members of the set from those that are not. For instance, CARs™ might search for the indifference surface that separates the priors that justify one strategy option over another.

Level Sets. We have also included derivations employing algorithms that efficiently find level sets of a function. Such algorithms are not nearly as thoroughly investigated as those that find extremal values. There are few algorithms designed to generate a maximally diverse set of points that satisfy some constraint. We experimented with variants of existing evolutionary search methods as promising candidates for generating members of level sets including some that have proved promising in our previous work. In order to localize the boundary of level sets in high-dimensional spaces, we extended known classification algorithms so that they will suggest which next sample points will provide the greatest leverage for further refining the emerging boundary.

Derived Contexts. A key motivation behind the design of our technology is that it enables facile manipulation of what we call compound computational experiments. These may be defined as a series of queries to the underlying model or data designed to answer a particular question of interest and are a key to implementing the inductive, RAP™ approach to robust decision making. These compound computational experiments are supported in CARs™ 2.0 through mechanisms we call derived contexts. These make it easy for users to construct the compound computational experiments needed for robust decision making in their particular application. A derived context is a particular repackaging of the inputs and outputs of one or more calls to a model useful for answering a set of questions. All derived contexts that were previously supported in CARs™ 1.0 have been evaluated for re-implemention in CARs™ 2.0. In addition, a number of new derivation methods have been constructed. These include:

- the ability to combine a search generator with a context to make one or more inputs in the foundation context outputs in the derived context, and possibly one output in the foundation context can become an input in the derived context;
- a variety of sensitivity analysis derived contexts, which run a generator on the foundation and then calculate sensitivity measures on inputs in the foundation context which are then outputs of the new derived context (both local and global sensitivity analyses are now supported);
- a "fitted model" capability that builds a derived context by fitting or smoothing the data available in the foundation context (currently available algorithms are linear fitted model, or n-nearest neighbor smoothing);
- regret derived context where a output is added to the derived context that calculates the "regret" of a chosen strategy for the given input scenario, relative to other strategies that might have been chosen;
- a generator context, that allows the nesting of generators, where, for example, a full factorial design on some inputs can be used to create a matrix of starting points and some hill climbing search generator started from each of these.

Some of these derivation methods have revealed the need for enhancements to CARs™ mechanisms for control and data flow, and have stimulated innovations that now are part of our foundation technology.

Scripting. A major innovation has been the introduction of Java scripting in CARs™. Previous versions of our technology had used special-purpose scripting language to store transcripts of GUI actions. We re-implemented this feature in CARs 2.0, replacing that special purpose language by Java itself. This allows scripts to be used for a much wider range of purposes than just saving GUI activities for later replay. All of the classes used in CARs™ internals can now be made available to users to customize the software environment at runtime. We can now use this feature to create customized GUI's for particular applications and for rapid prototyping. This has proven to be a unexpectedly powerful approach. As a consequence, the rapid prototyping interface, a major goal of this

project, now exists. We are also getting numerous features not anticipated to be among the goals of this effort "for free".

Graphical User Interface. Significant advances of CARs™ 1.0's graphical user interface (GUI) have been made. The main panel has been revised, and the logical hierarchy of panels and features redesigned. Multiple new graphical types have been added. A facility for "drilling down" by mousing on a graph to generate a case in a different context has been implemented, along with the ability to capture cases, or bookmarks by mouse actions on a visualization window. As described above, the new scripting facility includes a feature that allows a new main panel to be devised which will come up on CARs™ initiation, with each button tied to a CARs™ script that will execute a series of actions should the user invoke it by clicking on the associated button. This allows us to use scripts to customize the system's look, feel, and functionality for specific problems, models, or user community. Through the ability to cusomize by simply executing a script on system startup, we free ourselves from ever having to alter CARs™ code in order to proto-type some application of it, which greatly facilitates the use of our system for rapid prototyping. (Of course, it may be desirable on occasion to convert that prototype to a dedicated system through such reimplementation.)

A major new component of GUI support, the context browser, has also been implemented. This browser provides the user with a graphical summary of all contexts currently created in the CARs™ environment, and the relationships between them. This interface gives the user visibility of a wide range of status information, such as the number of records currently available in any context, storage management policies for each context, the inputs and outputs of each context, and so forth. Thus, the context browser is a major new tool for helping users manage the complexity of the computation environment they create in CARs™.

All of the technical goals established for this Phase II project have been fulfilled.

III. Benchmarking End-User Utilization and Rapid Prototyping

A major subcontractor during the course of this project was Thoughtlink, Inc., a firm specializing in the analysis of large models and simulations, evaluation of data visualization

and collaboration tools, developing and administering user surveys, and identifying requirements. What follows are the reported findings from Thoughtlink resulting from various user interface experiments and interview efforts. What follows is excerpted from "Exploring Rapid Customization of CARs™: Results from User Workshops and Interviews"1, the final report to Evolving Logic.

1. Project Background

This effort, which started in April 2001, supported Evolving Logic's Phase II Small Business Innovative Research (SBIR) project for the Defense Advanced Research Projects Agency (DARPA). ThoughtLink's primary tasks as originally stated in the Statement of Work with Evolving Logic included:

- Understand CARs™ and its applications
- Design experiments and data collection instruments
- Conduct workshops and user interviews
- Analysis and documentation of results

Over the course of this 2-year project, ThoughtLink had numerous face-to-face meetings and teleconferences with Evolving Logic to review CARs™, its applications, and various implementations. ThoughtLink installed and used both the Linux and Windows versions of CARs™, running different models, to increase their understanding of CARs™, its applications and the different visualizations produced by the software.

ThoughtLink conducted interviews and workshops with recent and potential users in order to gain insights into how well they were able to understand CARs™ and the visualizations it creates, and to understand user needs for rapid customization of CARs™ across various domains. Two user workshops were conducted: one in March 2002 and one in April 2003. At these workshops, users were given background information on CARs™, information on the model used, they were shown various output visualizations, and asked questions designed to assess their understanding of the visualizations.

The first workshop used the MESA/SM (Multiple Engagements of Strategic Arsenals with Stability Metrics) developed by Los Alamos National Laboratory (LANL). One of Evolving Logic's tasks for the Phase II SBIR was to deploy a robust planning module of

[1] Mukerjee, Kaushik and Julia Loughran, May 2003.

CARs™ at LANL. MESA/SM examines nuclear strategic force planning and stability in multi-polar engagements. It is used to help determine what combinations of strategic arsenals, target vulnerabilities, and national-level decision rules lead to stable configurations where first-strike advantages are minimized.

For the second workshop, CARs™ was used in conjunction with the Peer Competitor (PC) model. The Peer Competitor model is designed to evaluate the long-term policy problem of how to react to and interact with potential regional or global peer competitors using a theoretic model of peer-hegemon competition. Potential CARs™ users came from a population of students and professors from the University of California at Los Angeles (UCLA) campus and from RAND graduate students – all participants had an interest in peer-hegemon strategies.

Interviews were conducted with two divisions of Volvo. The Volvo Car Corporation (VCC) in Sweden had used CARs™ visualizations as part of their strategic planning for selecting new vehicle types. ThoughtLink later contacted Volvo's North America operation in Camarillo, CA to discuss their use of an updated version of CARs™.

In total, three different CARs™ user groups, totaling 20 people, were contacted in the context of gathering their feedback and understanding of CARs™ and its visualizations. Although this is a small sampling, the results from these interactions did provide insights into how CARs™ might be rapidly customized to fit a wide range of CARs™ users and applications.

The following sections of this report provide a brief overview of CARs™, the research methodology, results from the interactions with actual and potential CARs™ users, and recommendations for next steps.

2. CARs™ Overview

CARs™ is an exploratory modeling software package that allows users to explore problem spaces characterized by complexity and deep uncertainty. It is designed to assist in exploration across myriad alternative futures and examine the effects of possible actions contemplated by users. "CARs™ provides a general-purpose engine for reasoning with compound computational experiments by offering a convenient platform for linking the numerous functions necessary to support these experiments and user interactions. CARs™

can encapsulate a wide variety of sources for computational experiments such as simulation models, statistical analyses, neural net training processes, and use large numbers of such experiments to support reasoning and problem solving. The technology allows the capture of both knowledge about and the deep uncertainty regarding complex systems by representing infinite ensembles of alternative plausible models. CARs™ allows reasoning about the properties of such ensembles, and hence properties of ensembles of possible future events, through tailored constellations of computational experiments whose specification is guided by user goals and reasoning strategies.2"

CARs™ is designed to support Evolving Logic's "Robust Adaptive Planning™ or similar robust decision method. CARs™ utilizes either an existing model provided by the user or a new model generated from user input as a way to generate a number of scenarios. These results are displayed with CARs™ visualizations. The visualizations allow users to see patterns, understand the relationship between input factors, and pose new questions.

Traditional analysis tries to take away (or ignore) the complexity of a problem by forcing the analyst or model developer to provide answers to or values for unknown factors. Alternatively, the complexity remains when using CARs™ - and unknowns remain as hypothetical values versus foregone conclusions. This provides a much more realistic picture of the potential outcomes. Using CARs™, various strategies can be compared, robust strategies can be selected, and conclusions can be drawn, even under conditions of deep uncertainty.

Using CARs™, a variety of visualizations can be produced. These visualizations include, among others:

Landscape (or Box) Charts: These compare a variety of outputs and help array ensembles of related scenarios.

Region Plots: These are smoothed-out box plots that help to show trade-offs.

Line Charts: These charts provide the details of the phenomena that lead to the results displayed in any given square on a landscape chart.

3D landscapes: these landscape charts compare multiple dimensions of the problem space.

---

2 CARs™ description written by Steven Popper as part of the PC-CARs Workshop.

All of these visualizations are produced interactively. Users move slider bars to select attributes to be varied. In this way, they direct the way CARs™' interacts with a domain model. A variety of hypotheses can be generated and users see changing results.

Since CARs™ is customized for each client's problem domain, the knowledge elicitation step is an important part of Evolving Logic's overall process. They have created their own approach for knowledge elicitation called "XLRM." The XLRM process imposes coherence on the problem. The process is much appreciated by the user communities. It breaks the problem into a 2 X 2 grid. Each section of the grid is labeled with either X, L, R, or M. Each of these is described below:

X: These are the uncertainties outside of the user's control (e.g., future oil prices, competitor's pricing, etc.)

L: These are the levers, the actions that the users can take – a selection of particular levers aggregates into alternative strategies.

R: These relationships tie the other three areas together. These statements constitute the model.

M: These are the metrics and measures. Generally, there is no single best indicator – however there are often a series of key outputs of interest to the group commissioning the analysis.

The initial listing in the XRLM grid helps elicit more information for each of the categories, extracting all of the user's implicit assumptions. The XLRM grid is then used to help focus the users on what is most important to them; what they most want to do. This is an iterative process and often the diagram is circulated to other stakeholders. Finally, the information (when approved) is used to drive the model/scenario generator for the particular domain.

A sample XLRM for key factors to consider in MESA is shown below.

Sample XLRM for MESA/SM

Uncertainties Outside Control (X)	Levers Under Control (L)
• Formulations of "f" functions • Damage expectancy, p_K • Competitors' response strategies • Scenario: players, order of strikes • Weapon reliability • Defensive capability	• Own CF/CV strategies • EAs • Strategic force weapon level/mix • Defensive systems • Target weighting • Target coverage • Aggressiveness
Objectives and Measures (M)	**Relationships Between Factors (R)**
• Stability measures (suite) • FSS, Equilibrium • Robustness of strategy across scenarios; weapon reliability • Coverage goals over target assets • Risk of accident • Change in correlation of forces	• MESA/SM • National goals and forces, strategies, stability assessments • Own freedom of conventional action and relation to stability

Evolving Logic also asks for pictures of users' information, such as monthly reports, to discuss how these pictures are used to help decision makers. This is important for determining how in each use community people convince and are convinced by others. Rapid prototyping consists of crafting the CARs™ tool to yield views similar to these.

3. Research Methodology

The original plan for this task was to conduct workshops and user interviews with only one set of respondents: users of the MESA-in-CARs™ application. However, ThoughtLink felt it was important to get a wider range of views on the rapid customization of CARs™ since the software can support a myriad of applications dealing with deep uncertainty.

ThoughtLink learned CARs™ had been used successfully with the Volvo Car Corporation (VCC) in Sweden so they were contacted. Later, ThoughtLink contacted Volvo's North American operation. They are currently using a new application tool built upon CARs™. The MESA-in-CARs™ workshop was conducted and after that, another CARs™ application was selected for an additional workshop. Evolving Logic selected the Peer Competitor model. Both MESA and CARs™ are described in later sections of this report.

The goals of all the workshops and interviews were to:

Gain a better understanding of how different visualizations support potential CARs™ users in a variety of domains

Gather ideas for alternative visualizations

Improve the process of rapid prototyping of CARs™.

For each of the three user groups contacted, ThoughtLink developed questionnaires and worked with Evolving Logic to develop pre- and post-workshop questionnaires. The results from each of the three user groups follow. This report goes into the greatest level of detail for the PC-CARs™ (Peer Competitor) workshop; it was the most recent and had the largest group of participants.

4. Volvo Car Corporation Interviews

A. Volvo - Sweden

VCC in Sweden contracted with Evolving Logic to customize CARs™ so it could be used as a decision support tool in product planning for a new fleet of smaller vehicles. This CARs™ application takes into account a complex set of investment alternatives to show how they affect the rest of the manufacturing enterprise. For instance, it can help assist in addressng questions such as "When selecting a particular fleet of cars, what price do I need to charge to make a profit?"

In January 2002, ThoughtLink interviewed Charlotta Källbäck over the telephone to understand their experience with CARs™. Ms. Källbäck is the assistant to the VCC's chief executive officer on strategic product planning issues.

The CARs™ user community at VCC was drawn from members of the interdepartmental Product Planning team (which itself consisted of VCC staff from product planning, the project organization, and Ms. Källbäck representing the Office of the President). This group's interaction with CARs™ began with an organization-wide search for meaningful data that could be applied to CARs™. After receiving the relevant data from VCC, Evolving Logic then presented simplified simulation outputs in picture form (due to platform compatibility issues – CARs™ did not run on the VCC's information systems environment).

Ms. Källbäck offered numerous insights and suggestions on VCC's experience working with Evolving Logic to develop a customized version of CARs™ and on VCC's experience with CARs™. Of note are the following observations:

The process of exchanging data back and forth between VCC and Evolving Logic was time-consuming, partly due to the time difference between California and Sweden. This process could be shortened if the number of iterations were reduced. However, face time (when Evolving Logic came to VCC) was extremely valuable.

The interactive nature of CARs™ visualizations was well received. Participants wanted to see more and more visualizations.

One benefit of using the visualizations across a number of heterogeneous groups was that it helped in developing a common language and a common understanding of what factors were important for the different user communities. She said CARs™ helped users develop a common understanding of relevant issues, particularly the sensitivity of important parameters.

Regarding the usefulness of CARs™ in supporting VCC decision-making, Ms. Källbäck emphasized that the user group's interaction with CARs™ was a test of VCC's possible adoption of software to support its decision-making processes. Actual decisions may or may not have been different despite using CARs™, primarily because of a variety of factors external to the model that affected how (or whether) it was actually used.

ThoughtLink attempted to contact additional VCC contacts. This failed: either the contacts felt they did not have enough experience with CARs™ to comment or they were unable to be contacted due to scheduling difficulties. Therefore, Ms Källbäck was the only person interviewed from this user group.

B. Volvo – North America

Evolving Logic worked with Volvo's North American operation to develop a customized application of the CARs™ environment to support the company's decision-making process in the development of new products and services. The application is called RAPNOW, for Robust Adaptive Planning™ for New Ownership Models.

The information here is based on a phone interview conducted by ThoughtLink on May 16, 2003 with Mr. Benny Sommerfeld, project leader at the Volvo Monitoring and Concept Center in Camarillo, CA.

Mr. Sommerfeld's work with Evolving Logic dates back to 1998 and includes a variety of special projects as well as the RAPNOW application. The RAPNOW application runs onsite at Volvo's Monitoring and Concept Center on a Linux machine.

Mr. Sommerfeld and his team have used RAPNOW to answer unique questions to analyze the costs and benefits of various strategic options or "business cases." Volvo offers cars and services and they use RAPNOW to look at new customer offers consisting of various services and vehicle offerings. Their primary focus area is on environmental factors. The strength of RAPNOW is its allowance for uncertainty throughout the analytical process.

Mr. Sommerfeld's group has used RAPNOW to conduct investigations on issues of strategic intent, to better understand whether certain strategic options are profitable or nor profitable, to experiment with sensitivity analysis, and to look at the robustness of various options. To make a business model work, Volvo needs to look at factors that cannot be controlled. RAPNOW allowed Mr. Sommerfeld to do this, and showed the potential of different strategic options. RAPNOW also helped formulate which activity plans had to be carried out to achieve strategic goals.

One significant application area for RAPNOW was to examine new ways of structuring Volvo's supply chain relationships among Volvo (as an OEM – original equipment manufacturer), its suppliers and vendors, and its customers. Volvo had been looking at these relationships and applied RAPNOW to explore the robustness of various strategies and options. RAPNOW emphasized the importance of looking at Volvo's strategic suppliers. For example, "end of life" legislation in Europe requires suppliers to take back unused inventory from automobile manufacturers. Volvo had been looking at this issue from a life-cycle analysis standpoint, to meet legislative requirements and to be a good corporate citizen. By changing the way they work with suppliers, Volvo realized they would both gain revenues over the life cycle of a product (specific vehicle model), yet without increasing the cost of ownership to the end customer. This realization was not a direct outcome of RAPNOW. Volvo knew beforehand that they wanted to pursue this

strategy. However, the value from RAPNOW was what it told Volvo about the robustness of this particular strategy – coordinating with strategic suppliers to share costs would result in increased revenues for both, at no change in cost to Volvo customers.

Other general comments from Mr. Sommerfeld were that RAPNOW highlighted which parameters are important to look at when you have already decided on an appropriate plan or strategy. It is then possible to go through a "reverse engineering" process to find out which numerical values will help you to achieve your goals. He emphasized that CARs™ is a rich and complex tool, and that he and his group benefited the most from having Evolving Logic available to make sense of the environment and the modes of interaction with it.

Of the multitude of variables available in RAPNOW, Mr. Sommerfeld stated that currently he is only looking at a few that make sense for the specific analysis he is conducting. Partly due to time constraints, he and his group have only "scratched the surface" of CARs™' capabilities, and he hopes to delve deeper into additional features and functionalities. Mr. Sommerfeld would like to have a more "crisp and clear" interaction with the model, and to this end, Evolving Logic should consider his recommendations:

RAPNOW needs a more user-friendly interface. Axis variables need greater clarity, and it should be easier to read the legends.

A user manual is very important to derive the most value from RAPNOW. The current user manual is in its infancy, and it still takes a lot of time to work within the CARs™ environment.

Certain RAPNOW parameters should 'pop' out more. For instance, it should be possible to go to a chart and click on a certain area for further explorations.

Evolving Logic should consider incorporating different levels of interface and analysis in RAPNOW, tailored to the user's level of expertise, which would allow them to move to deeper levels of analysis as their experience and comfort level with RAPNOW increases.

When asked if he has had the opportunity to present RAPNOW outputs to other groups at Volvo, Mr. Sommerfeld stated that he uses RAPNOW to tell a story to different audiences, and for this purpose, RAPNOW charts give a good representation. Although

sporadic interactions has made it challenging to work with RAPNOW, Mr. Sommerfeld attests to the value of the tool, and plans to spend more time with it.

RAPNOW has met Mr. Sommerfeld's expectations, and he stated that access to Evolving Logic has been very positive and helpful. For instance, the latest version of RAPNOW has more Help information and the new toolbox has been helpful.

He added that Evolving Logic's 'XLRM' methodology has made a significant difference in his team's work. As an extension of XLRM, RAPNOW has been very valuable in generating questions needed to conduct an analytical process. Without RAPNOW, Mr. Sommerfeld said that the GIGO effect (garbage in-garbage out) would be apparent. RAPNOW helped him to focus on working with precise and high quality data inputs, to ensure that data outputs were useful.

Mr. Sommerfeld strongly suggested that Evolving Logic consider offering more training in its core XLRM methodology, since using the RAPNOW tool alone, without the XLRM framework, would not deliver as much value to CARs™ users. RAPNOW could be viewed as a proof-of-concept of the XLRM methodology.

5. MESA-in-CARs™ Workshop

The MESA-in-CARs™ Workshop was held at the same time as the Fourth Nuclear Stability Roundtable in March 2002. This provided the opportunity to gather a number of potential MESA/SM users. The workshop was held at the Defense Threat Reduction Agency (DTRA) in Springfield, VA. For several months prior to the workshop, Evolving Logic and LANL personnel had been working to integrate the MESA model with CARs™. LANL wished to explore whether CARs™' visualizations could provide new insights about model results that had not previously been obtained from MESA's numeric tabular output.

The potential user group attending the workshop included seven people – three from DTRA, two from LANL, one from US Strategic Command (STRATCOM) and one from the Institute for Defense Analyses. They all had experience in the domain of strategic stability. Their roles ranged from analysts/researchers to program managers/project leaders. They also varied in their experience with the MESA/SM model – three had actually used the model and three had seen it. Only one person was completely unfamiliar with the model.

The attendees self-assessed their familiarity with MESA/SM on a scale of 1-5 (where 1=Not Familiar and 5=Very Familiar). Their average score was 3.64 (mean of 3.5).

One purpose of the workshop was to obtain user feedback on the overall concept of robust decision methods and in particular Robust Adaptive Planning™ using CARs™. Traditional defense simulation typically requires assigning single values to inputs, even if those values are not known in the real world, and the outputs represent single measures. In contrast, CARs™ provides ensembles of alternative scenarios or model specifications based upon a range of possible inputs.

Another purpose of the workshop was to assess the users' understanding of different types of output visualizations.

The workshop presented an overview of MESA/SM, an overview of CARs™, and a presentation of CARs™ output that showed a relevant storyline. The different output displays shown included Landscape Box Plots showing landscapes of alternatives defined by vectors of alternative model inputs and Line Charts showing the time paths that led to these outcomes. Eighty percent (4 of 5) of the respondents felt the Landscape Box Plots were easy to understand but only 40% (2 of 5) felt the line graphs were easy to understand.

One participant felt MESA-in-CARs™ would be more useful as an analysis tool vs. a tool to present results, however, another participant felt the CARs™ output would be useful in a presentation and he stressed the importance of being able to copy the CARs™ output into a PowerPoint presentation. Currently this can be done by saving the output displays as graphical images (e.g., in a .jpg format).

Workshop participants had differing opinions on which strategic engagement simulation was the most accurate and were skeptical of results derived from other simulations, including MESA-in-CARs™. ThoughtLink observed that what may have been more useful would have been to tie CARs™ to all of the models in this domain so that the visualizations from all of the models could be explored within a consistent framework and the implications of each of different assumptions could then be examined in a search for strategies that would be robust across alternative model specifications as well as alternative assumptions about variable values. As indicated by VCC's experience, CARs™'

visualizations are a powerful tool to support a rhetorical process that helps people understand differing views and come to consensus on key factors in the problem space.

Participants offered many suggestions and critiques for enhancing MESA-in-CARs™ to make it easier for the user community to understand and adapt CARs™ to their requirements. The following are some highlights:

Provide Right Level of Detail: CARs™ outputs need to be crafted to the right level of detail to meet user requirements. Depending on the user – specifically, key decision makers at DTRA, STRATCOM and other organizations – MESA-in-CARS™ may be delivering more detail than they are interested in seeing, or too little.

Integrate Visualizations with Microsoft Office Products: MESA-in-CARS™ should be viewed as an analytical tool to identify regions of interest, and then to generate graphs which can be inserted into commercial off-the-shelf programs (like Word or Excel) for presentation to decision makers. CARs™ should be used to support the analytical process.

Provide Access to Model Logic: CARs™ should allow easier access to underlying assumptions that are driving the model. At the MESA-in-CARs™ workshop, many users wanted to understand the MESA models underlying logic.

Finally, the MESA-in-CARs™ workshop yielded some useful insights on presentation modes and structuring the workshop agenda. ThoughtLink's observations include:

Allocate More Time for Feedback: More time needs to be allocated to eliciting information and feedback from workshop participants. Another approach to structuring the workshop might be to have users discuss their backgrounds, any questions they might have, and the models/tools they currently use. CARs™ (or MESA-in-CARs™) should then be presented after gaining a sense of who the users are.

Test Users Understanding of CARs™ Concept: Users may have missed the important point that CARs™ is not a visualization interface that 'sits' on top of an underlying model. Evolving Logic and ThoughtLink could develop pre- and post-tests to present this concept to users and make sure that it is absorbed.

6. PC-CARs™ (Peer Competitor) Workshop at UCLA

A. Workshop Overview

In support of Evolving Logic's modeling of the analytical framework presented in The Emergence of Peer Competitors published by RAND scholars, a CARs™ workshop was held in late April 2003. This workshop is referred to as PC-CARS™.

The PC-CARs™ workshop was designed to explore how leading-edge computational tools might support national security policy decision makers and gather feedback on the rapid customization of CARs™. CARs™ was presented as an analytical framework for understanding the long-term policy problem of managing interactions with potential regional or global peer competitors. The workshop intended to provide participants with an introduction to and a practical demonstration of the CARs™ environment, illustrating the possibilities of connecting analytic research and policy work through current computing technologies.

This workshop used a different methodology from the earlier MESA-SM workshop: users were divided into two groups and each group was shown the same visualizations but in a different order (alternating between showing landscape or line chart visualizations first). This was intended to provide more information about the correct order to present the different visualizations in order to increase participant understanding.

Originally intended for the student and faculty community at the University of Southern California, the workshop was shifted to UCLA's School of Public Policy with the cooperation of UCLA professor Dr. Michael Intriligator. Invitations were made to four groups:

UCLA undergraduate students enrolled in Dr. Intriligator's 'Nuclear Strategy' honors course

UCLA graduate students enrolled in Master's and Ph.D. programs in economics, political science, and policy studies

UCLA faculty members

RAND graduate school students and alumni

ThoughtLink developed online registration on the Evolving Logic corporate website and prepared the workshop's read-ahead package provided to the participants. See Appendix A for a copy of the read-ahead document.

The PC-CARs™ workshop was conducted in 2 sessions on the UCLA campus on April 28, 2003. The 2-hour workshops were held from 1-3 p.m. and 7-9 p.m. The workshop was designed to be interactive – time was allotted on the agenda for participants' questions and discussion. Each 2-hour session consisted of a 30-minute (roughly) background discussion led by Steven Popper of Evolving Logic, and an hour-long discussion of the PC-CARs™ model led by Rob Lempert of Evolving Logic with approximately 30 minutes for discussion. At several points during the workshop, participants were asked to respond to specific questions on questionnaires that were handed out prior to the workshop.

The background portion of the workshop covered the following topics:

Introductions

Issues / Peer Competitor model

How does lack of perfect information affect outcomes?

Peer strategies have implications for its growth in power, relative to hegemon power

Participants respond to PC model questions

Discussion of PC model questions

CARs™ approach

Questions

The agenda for the PC-CARs™ model presentation portion of the workshop was as follows:

Presentation of CARs™ Hegemon visualizations

Presentation of CARs™ Challenger visualizations

Questions on Challenger views

Discussion

Finish questionnaires

The PC-CARs™ visualizations were presented in two different ways to test participants' understanding of concepts presented in the background portion of the workshop. In the 1:00 PM session, participants were first shown regret landscape views for the range of possible Hegemon strategies, followed by line chart views displaying the Challenger's base power growth rate. In the 7:00 PM session, the presentation order was reversed – participants first viewed the line charts (based on the Hegemon's chosen strategies) and then saw the regret landscapes. In both sessions, after viewing the landscapes and line charts for the Hegemon, participants were then asked to evaluate which strategies were optimal for the Challenger, based on a different set of Challenger-driven regret landscapes and line charts.

Participants were asked to respond to two questionnaires. The first was a pre-workshop questionnaire to collect basic demographic information and assess the individual's familiarity with personal computers and software, and experience in the modeling and nuclear strategy decision-making domains. Answers to the pre-workshop questionnaire were gathered either over the phone during the registration process or filled out on paper by the participant prior to the workshop. The second questionnaire consisted of a blend of quantitative and qualitative questions to gauge participants' understanding of CARs™ concepts and visualization outputs. See Appendix B and Appendix C for copies of the pre-workshop and workshop questionnaires. As indicated above, rather than attempt to capture how well the participants absorbed these concepts at the end of each session, ThoughtLink asked participants to respond to the relevant questions on the PC-CARs™ model and the landscape/line chart interpretations immediately after they were presented. These questions were later graded by Evolving Logic.

B. Participants

The group of PC-CARs™ workshop participants represented a cross-section of the four invited groups mentioned earlier. Although 15 people had signed up for the workshops, 11 people actually attended.

The pre-workshop questionnaire shows the following data on the respondents' education, computer skills, familiarity with computer modeling software in any domain, and familiarity with nuclear strategy decision-making simulations in particular:

- Highest level of education: 36% high school diploma (4 respondents), 27% PhDs (3 respondents), 18% Master's degree (2 respondents), and 18% other degree (2 respondents)
- Level of computer skill: 4.91 mean (7-point scale), 5.0 median; suggesting somewhat above intermediate skills
- Most-cited software programs that respondents use on a regular basis include Microsoft Office (Word, Excel, Internet Explorer), and statistical software (Analytica, Stata, SAS)
- Expertise in computer simulation/modeling software applications, any domain: 3.27 mean, 3.14 median; suggesting somewhat below intermediate-level expertise
- Expertise in computer simulation/modeling software applications, in nuclear-strategy decision-making domain: 1.55 mean, 1.0 median. This suggests somewhat above beginner-level expertise.

Group	Participants, 1:00 PM Session	Participants, 7:00 PM Session	Total
UCLA undergraduate students enrolled in 'Nuclear Strategy' course	2	3	5
UCLA graduate students	0	1	1
UCLA faculty members	2	0	2
RAND Graduate School students and alumni	3	0	3
	7	4	11

Table 1. Breakout of PC-CARs™ Workshop Participants

C. Domain Understanding of PC-CARS™

To test respondents' understanding of the concepts of the Peer Competitor analytical framework modeled in PC-CARs™, an in-workshop questionnaire was administered to present pre- and post-test questions to measure how well the concepts were absorbed. Following a presentation of the PC framework, participants were asked the 'pre-test' question: "What factors are most important to determining the best choice of strategy?". The CARs™ approach was then presented, followed by the PC-in-CARs™ environment - the actual model. After presenting the model, participants were asked the post-test question: "Now that you have seen the Peer Competitor model - what factors are most important for determining the best choice of strategy?" to assess how much their understanding of the PC framework changed as a result of viewing a representation of this framework in the CARs™ environment.

The anonymous responses to these questions were compiled by ThoughtLink and presented to Evolving Logic, who ranked the responses on a 5-point scale (based on the following definitions:

- 5 - respondent understands concept very well, could understand other concepts and visualizations equally well
- 4 - intermediate score
- 3 - respondent has a basic grasp of the concept, but may struggle with interpretations of other concepts or visualizations if presented
- 2 - intermediate score
- 1 - respondent has limited or no understanding of concept/visualization that was presented.

This ranking system provided a basis for understanding whether the CARs™ environment itself played a role in helping participants to understand the Peer Competitor model.

The participants' responses to each question had two rankings, one from Steven Popper (SP) and one from Rob Lempert (RL). The rankings for all 11 respondents were combined to yield a mean score and standard deviation. Results of the rankings are as follows:

- For the pre-test question (before presentation of PC-CARs™), the mean scores are (3.00 (SP), 3.78 (RL)), and standard deviation is (1.15 (SP), 0.79 (RL)).

- For the post-test question (after presentation of PC-CARs™), the mean scores are (3.88 (SP), 3.56 (RL)), and standard deviation is (1.35 (SP), 1.07 (RL)).

Steven Popper and Robert Lempert's rankings of participants' responses are slightly divergent. While Mr. Popper's mean score increased from 3.00 to 3.88 on the 5-point ranking scale (a +0.88 difference), Mr. Lempert's mean score slightly decreased from 3.78 to 3.56 (a -0.22 difference). Evolving Logic's differing assessments of participants' understanding of the PC model indicates that there may be opportunities in future user interactions to better define what constitutes 'understanding' of the CARs™ environment. In addition, the small size of the participant group (11 total) adds to the difficulty of drawing firm conclusions from the numerical rankings obtained from the questionnaires. Larger CARs™ user groups would yield more quantitative and qualitative response data to arrive at deeper observations and meaningful conclusions.

D. Results from Visualization Questions

To experiment with how different modes of presenting PC-CARs™ visualizations affect user understanding, Evolving Logic and ThoughtLink used different presentation orders for the 1:00 PM and 7:00 PM workshop sessions. In both sessions, participants were presented with visualizations generated from the same set of scenarios in which assumptions regarding a Hegemon's and Peer Competitor's initial attributes and the PC's initial strategy were known. However, in the 1:00 PM session, Evolving Logic first presented the strategic space as seen in the regret landscape charts, followed by line charts displaying the PC's base power growth rate. That is, this presented the high level overview focusing on the examination of the robustness properties of alternative strategies. It was then followed by a drill-down examination of time paths showing the underlying behavior that led to differing results to alternative strategies across different specific scenarios. In the later workshop session, the presentation order was reversed: line charts were presented first, followed by regret landscapes – a "bottoms-up" approach. See Appendix D for screen captures of the charts that were presented.

Participants were asked to write their interpretation (narrative) of the regret landscape chart and these written interpretations were later graded by Evolving Logic on a 1-5 scale. The mean scores are (3.25 (SP), 4.10 (RL)), and standard deviation is (1.54 (SP), 1.22 (RL)),

suggesting average to above average interpretations of the regret landscapes, with a slightly wider distribution of scores for one Evolving Logic evaluator.

Note: Responses to the line chart interpretation question were not ranked, based on a decision after the workshop that the wording of this question may have confused respondents.

Based on responses to the 'Visualizations Questions' section of the workshop questionnaire, respondents appear to have understood the visual outputs of PC-CARs™ "somewhat well." The line plot, 2-D landscape and 3-D landscape interpretations ranked 4.0, 5.18 and 4.82 respectively (4 being the middle score in the 7-point ranking scale that was used), suggesting that respondents had more difficulty interpreting the line chart of peer base power growth rates, compared to the regret landscapes. This agrees with similar findings from the MESA-in-CARs™ workshop where line charts were also less well understood. Respondents had much greater familiarity with the line chart display format (most responded that they had seen this type of display before) than with the 2-D and 3-D landscape displays (about 80% had seen neither type of visualization before). It is difficult to explain why the line charts may be more difficult to understand since people seem to have greater familiarity with them.

The use of color to display different ranges of value was not a problem for respondents, with the exception of one respondent who is color-blind (but did not indicate to which colors). The majority of respondents were able to infer a variety of conclusions from the CARs™ visualizations – these inferences were not graded – but can be reviewed in Appendix E (Questionnaire Responses).

Respondents had many suggestions when asked in open-ended questions for additional features and functionalities they would have liked to have seen in CARs™' visualizations, and for other types of visualizations such as the line charts and landscape views they saw. Overall, respondents reacted positively to the 2D and 3D regret landscapes, and seemed to have the most trouble comprehending the line charts displaying the PC's base power growth rates. Participants also had some difficulty in following the definitions of Hegemon and PC player attributes (based on the triplet letters such YNY, YYY etc.). Suggestions included:

- Suppress regret values for "cells" not significant in reaching final conclusions (e.g., reduce it to a comparison of "minimal" vs. "high" regret)
- Provide explanations of Hegemon and PC attributes (YYY)
- Focus on landscapes; de-emphasize line plots of base power growth
- View multiple line plots at the same time
- On the 3D landscape, show greater color gradation using more colors, rather than showing different layers and boxes

In terms of additional visualizations, participants offered a few (but original) suggestions. A sampling of these comments:

- View a 4x4 matrix that responds to slider bars as much as the regret landscapes
- Rotate the 3D landscape – to take advantage of the fact that it is 3D
- Interact with visualizations like the cockpit of a flight simulator, to allow user to learn from experience
- Compare both Hegemon's and PC's choices and regrets on the same landscapes and line plots

Comparing participants' understanding of the Hegemon and Peer Competitor strategies based on the presentation order of CARs™ visualizations (the 1:00 PM session viewed the regret landscapes followed by line charts of the PC's base power growth rate; the 7:00 PM session had the order reversed) the following results were obtained:

- The 1:00 group scored their understanding of the Hegemon strategies at 4.29, and the Peer Competitor strategies at 4.57. (One respondent ranked their understanding at '1' – the lowest score, meaning poor understanding – for both strategies.)
- The 7:00 PM group scored 5.67 for their understanding of both the Hegemon and PC strategies. (This is based on one non-response from a group of 4 respondents.)

These results indicate that the 7:00 PM group had a better understanding of CARs™ visualizations and the various strategies that were represented when they were presented

with the line charts followed by the regret landscape views (keeping in mind a smaller group of respondents and one non-respondent in the 7:00 PM session). The 7:00 PM group's score was a full point higher than the 1:00 PM group's – the 5.67 score placed the 7:00 PM group's understanding at an intermediate/high level, while the 1:00 PM group's scores (at 4.29 and 4.57) were closer to the 'Somewhat' midrange point on the scale. The implication here is that differing order of presentation of CARs™ outputs might have an impact on respondents' understanding of the PC model and the range of strategies available to the Hegemon and Peer Competitor. The results also support the findings presented previously in this report that workshop participants had greater difficulty in interpreting the line chart displays compared to the regret landscapes. There may be an opportunity in future workshops to ask pre- and post-test questions specifically measuring respondents' conceptual understanding of the varying strategies in the PC model. Finally, a slight majority of respondents (63%, 2 non-responses) were comfortable with the presentation order that was used in the workshop, and 2 people offered suggestions in this area, including incorporating more current examples in the model, and tailoring the presentation order to the experience level of the workshop participants.

E. Relationship Between Background Experience and Understanding

The workshop questionnaire asked specific questions about respondents' experience in a variety of domains. In general, the overall group had moderately low experience with PC-like models (2.82 on a 7-point scale) and data visualization tools like CARs™ (2.91). Respondents ranked their experience in the nuclear strategy decision-making domain, however, somewhat higher at 4.45.

Most of the participants (72%) did not visit the Evolving Logic website prior to the workshop, opting instead to review the read-ahead materials for background information (91%). About 73% of the group found the read-ahead information useful in understanding the PC model and the CARs™ environment.

Interestingly, the group had a neutral response (4.2) when asked about the usefulness of the workshop in teaching national security policy decision-making, and an equally neutral reaction (4.33) on the adequacy of the PC model. Although the group found the PC-CARs™ workshop to be somewhat educational and informative (5.22, one level above neutral), there was little change in their self-assessment of how much their expertise in

national security policy decision-making changed (4.45 before the workshop, 4.55 after the workshop).

F. additional PC-CARs™ Workshop Results

The PC-CARs™ workshop at UCLA provided direct interaction with potential users of the model, each of who had varying degrees of experience and expertise within the nuclear strategy decision-making domain. Presenting PC-CARs™ to this group, and collecting their input and assessment of the model, provided valuable insights on the potential acceptance and adoption of PC-CARs™ by other audiences. Some additional workshop findings include:

- Consider User Groups Domain Expertise for Future Workshops: The overall group's unchanged assessment of its expertise in strategic decision-making, before and after presentation of PC-CARs™, shows that factors specific to the CARs™ environment, and varying modes and orders of presentation of the model, are important factors to consider when planning future user interactions with PC-CARs™. This is particularly important when considering that the UCLA group ranked itself on the pre-workshop questionnaire as having somewhat below intermediate-level expertise in computer simulation/modeling software applications in any domain, and somewhat above beginner-level expertise in domain-specific software applications within nuclear-strategy decision-making. The group's self-assessment of its domain and non-domain knowledge of simulation/modeling software applications indicates that beginner-level user communities may derive greater value from PC-CARs™ workshops when these interactions are more targeted to their needs and concerns.
- Need for Legends on Visualizations: The group's responses indicated a high degree of understanding of the Hegemon's and PC's strategic options and the implications of choosing one course of action over another. Yet, many of the quantitative ranking questions and qualitative questions on their understanding of the visualizations show that respondents had difficulty in understanding aspects of the visual displays that were presented. Specifically, there are

recurring comments stating that the users had problems interpreting the line chart displays of PC base power growth rates, and the attribute definitions (YYY etc.) displayed on both the line charts and regret landscapes.

7. Findings and Next Steps

The PC-CARs™ workshop generated measurable results on respondents' understanding of the CARs™ environment within the specific domain of nuclear strategy decision-making. These results can help Evolving Logic to continue development and refinement of the PC-CARs™ model for possible future interactions within similar user communities, mostly within the academic and research communities. It is conceivable that different versions of PC-CARs™ could be generated, based on the results of this workshop, that are customized to meet the requirements of specific audiences. For instance, a version of PC-CARs™ catering to the decision support and research needs of professional researchers and analysts of nuclear strategy conflict issues could be explored.

It is also conceivable to build on the UCLA workshop and target PC-CARs™ as a classroom learning and research application for undergraduate and graduate students of political science, national security policy, economics and other disciplines.

Based on the multiple interactions that Evolving Logic and ThoughtLink have had during the course of this work with Volvo, MESA, and other groups, the following steps could be pursued in the future to boost user acceptance of the CARs™ environment and enable more rapid development of CARs™ models in other domains:

Expand Knowledge Elicitation Methods: Evolving Logic's XLRM process described earlier has proven to be quite successful for knowledge elicitation and users have commented how it has benefited their understanding of CARs™. Some suggested that a better understanding of this process helped them better understand the software. In addition to expanding the XLRM process, Evolving Logic should consider conducting the process in a distributed environment using collaboration technologies. So far, this interactive process has taken place face-to-face. In addition to the cost savings associated with reducing travel expenses, developing a persistent, collaborative environment might help provide input that participants may not think of during a time-limited face-to-face meeting. A web-based

asynchronous environment would be available for them to add information as they think of it.

Consideration of Presentation Order to Facilitate Understanding: The findings from the PC-CARs™ workshop indicate that understanding is facilitated if the line chart graphs are presented first, followed by the landscape charts. This finding is based on a small sample size, so additional research is needed, but these initial findings indicate that the line charts should be presented first.

Different Visualizations for Different Users: An important finding is that CARs™ needs to produce different types of visualizations – those to be used for analysts and those to be used by higher-level decision makers. In addition, Evolving Logic may wish to consider the benefits of organizing the CARs™ environment – or the user interface – to cater to users' needs based on their level of expertise within a particular domain, their job function, or other factors. Allowing the user to self-select their mode of interaction with CARs™ (beginner, intermediate, advanced, for instance) may minimize confusion and boost acceptance as users' familiarity and expertise grows.

The analysts interviewed from both VCC and the MESA-in-CARs™ workshop indicated data and model output must be massaged so it makes sense to the decision maker. It is important that the results "tell a story" and make sense.

A number of workshop attendees brought up the fact that CARs™ output needs to be crafted to the right level of detail for the potential audience. More information is needed to identify which types of charts provide the greatest level of insight for the spectrum of potential users. In the MESA-in-CARs™ workshop, one user commented that the Region Plots would be much more useful to policy makers. Another suggested the visualization should match the organizational culture that will be viewing the chart, e.g., the Department of Defense's affinity for "stop light" charts. The 3D charts shown during workshops were thought to be a bit too complex for policy makers. Instead, it was suggested that these charts might be more useful to an analyst who could understand the complexity shown in the display. The participant that made this remark said, "As an analytic geek, I love that chart, but I would never show it to my boss. He'd puke on it."

Adding Collaboration Capabilities: One benefit of CARs™ is that it presents visualizations that can be used to develop a common understanding or to argue different viewpoints. One significant improvement for CARs™ would be if it could run in a distributed environment – potentially a web-based intranet with corresponding collaboration tools (e.g., video teleconferencing or text chat) so interested parties could interact and share their interpretations of the results. Adding the capability to annotate the visualizations would also be a useful feature.

As indicated in a case study prepared by ThoughtLink on software development methodologies, Evolving Logic could consider incorporating structured communications and collaboration technologies (particularly web-based learning tools and user communities) to streamline the CARs™ development process. This is particularly relevant if Evolving Logic pursues longer-term interactions with CARs™ users as mentioned above – online, asynchronous model development approaches combined with face-to-face interactions such as CARs™ workshops may yield methods and processes that enable rapid CARs™ prototype development, and lead to shorter time frames for development of finished 'products'.

Longer-term, Continuous Interactions with Users: Interactions with various CARs™ user communities have yielded many useful insights, but have lacked the continuous interaction to generate valuable user feedback. This information would be useful in the continuous development and refinement of CARs™. This is especially relevant if Evolving Logic seeks to maximize user acceptance and adoption of the CARs™ environment among user communities of decision-makers and subject-matter experts. Establishing a CARs™ user community and tracking their uses of the software could offer tremendous insights into potential changes to increase usability.

Adapt User Interface to be More User Friendly: Although UCLA and MESA participants and Volvo users could grasp the structure and dynamics of the CARs™ environment, a recurring theme that emerges from these disparate groups is the unwieldy interface and a need for a crisper, cleaner mode of interaction between users and CARs™. Future user interactions, especially those that engage users over the long-term as mentioned above, could focus on gathering as much user interaction feedback as possible to develop custom interfaces that meet specific user requirements. Users need to be able to easily create

CARs™ visualizations, understand the parameters used to create those visualizations, and easily compare visualizations with different parameters.

Develop User Manual and On-line Help: Evolving Logic has recently delivered CARs™ to Volvo in Camarillo, CA. The software has limited on-line help and lacks a robust user manual. In order to continue on the path of commercialization of this software, Evolving Logic needs to invest in materials to support users in the field. Increasing user adoption and acceptance of CARs™ depends to a large degree on users' ability to interact with the environment, and extract value from it, in the absence of Evolving Logic experts. Thorough and easy-to-follow documentation of the CARs™ environment will provide users the guidance they need when interacting with CARs™ outside of a structured workshop setting. It will also provide users with answers to common questions and problems they will experience (troubleshooting and Frequently Asked Questions (FAQs)).

Consider User Recommended Enhancements: Some users suggested enhancements to the user interface and visualizations. These include:

- Adding markers on new graphical outputs where the colors used to be or somehow highlighting where the colors changed when the input variables changed
- Add an inset visualization of the base case so visualizations can be compared

Add a capability to stack visualizations and make them opaque to see where the colors have changed.

Issue of Providing Models Underlying Assumptions: It was apparent from meeting with potential CARs™ users that it is very important that users understand the underlying model and its assumptions. It would be useful if users could drill down to the part of the model that is being activated to produce the results shown in CARs™. Users indicated that it was very important to them to understand why certain results are appearing. Most of the comments from users focused on the model vs. the actual visualization. There is a desire to access the model's assumptions from the visualizations produced by the model. Since it seems apparent that these two things (model and visualization) cannot be separated in the eyes of the CARs™ participants, Evolving Logic needs to find ways in which CARs™ can give users a better understanding of the model itself. These findings are helpful to consider as the product continues to evolve and improve.

IV. PLANS FOR SBIR PHASE III

The Phase III period of this SBIR project is geared to see the commercial application of CARs™ 2.0 and its successors. The initial plan for commercialization has three components. The first is to used CARs™ to develop custom applications geared to answer specific client concerns. These specific applications may be used to develop materials that will help clients better understand the nature of the environment in which they operate and to undeerstand and convey to others the options available for crafting robust strategies for meeting these challenges.

The send component builds upon the first. This consists of licensing of custom-built software systems designed on the basis of CARs™ 2.0 technology. The user interface front ends of these custom applications are tailored to lead specific client groups through a Robust Adaptive Planning™ process on their own to answer recurring problems in planning and strategy development and design.

The third component will, in the context of the first two, be to develop general applications built on the core technology chassis to produce software that may be widely sold either within well-identified market niches or to address specific problems widespread across business and/or government.

It appears that transition to Phase III will be seamless with the activities conducted under Phase II. Already, commercial engagements have been booked or are in prospect for the first two of the three types of engagement listed above.

Sales of CARs™-generated output:

- Volvo Car Company [four projects completed];
- Los Alamos National Laboratory [completed];
- International Sematech (research consotium of major semi-conductor manufacturers) [at inception];
- National Defense University Center for Technology & National Security Policy [at inception]

Licensing of CARs™-based applications software:

- Ford Motor Company [delivery of CARAFE system completed];
- Volvo Car Corporation [delivery of RAPNOW system completed];
- Genoa II Program of DARPA's Terrorism Information Awareness Office [development of RAPSAW system in process];

In addition to these efforts, several others are at the preliminary stage of discussion and therefore are only prospective at this point. However, we anticipate the pace and volume of these efforts to increase.

We are also engaged in serious discussions and planning efforts with a potential partner in a venture to bring out the first third stage product, an application for an industrial niche to be widely sold within that industry. Given the delicate early stages of this discussion and the proprietary concerns of both Evolving Logic and the potential partner, we will not discuss this effort further at this point. However, this has provided a useful indication that our vision for Phase III commercialization along the full range of activities listed above is a prospect with good chance of success.